In the preferred embodiment, a "node" is physically packaged as a single electronic circuit card and has multiple chips mounted thereon. However, a node could have different physical embodiments, e.g., as multiple circuit cards in a mother-daughter arrangement, or multiple cards in a special packaging arrangement, or even possibly as a single chip. The significance of the node is that communication within a node is easier than between nodes. A single circuit card was chosen as the node boundary because communication across a card boundary requires connectors, drivers and other hardware to support it, making it desirable to minimize the number of lines crossing the card boundary

In the preferred embodiment, certain functions are combined on a single chip or divided among pairs of chips. However, the exact number of chips used to perform functions may vary, and as described herein, devices are grouped by logical function rather than packaging. For example, it may be possible to combine the two DSW chips on a single chip, or even to combing the ARP and DSW chips on a single chip. It may alternatively be possible to implement the ARP as multiple chips. The logic functions required of the ARP, DSW and ASW are far less complex than a single processor chip, but multiple chips are used chiefly to obtain the requisite number of I/O pins.

In the preferred embodiment, the bus links are configured hierarchically, so that a local repeater services multiple local links. However, it would alternatively be possible for multiple input links to go directly from devices to the central repeater, and for the central repeater's output to go directly to all devices, without intermediate local repeater units. Many other variations in the topology or types of links are also possible.

In the preferred embodiment, the input links are half the width of the broadcast bus, requiring two bus cycles to transmit a complete command. However, the input links could be ⅓, or ¼ or some other fraction of the broadcast bus. The purpose of the multiple input links is to utilize the broadcast bus to the fullest extent possible. As the number of devices in the network grows larger, and the number of input links grows correspondingly, it may be predicted that narrower input links are still sufficient to fill the broadcast bus.

In the preferred embodiment, the local request links running from devices to ARP, and the remote request bus links running from ARP to the central repeater, are the same width. However, it would be possible to employ differing widths. For example, the local request links could be ¼ bus width while the remote request links are ½ bus width. Alternatively, the remote request link(s) might be full width, while only the local request links are less than full width.

In the preferred embodiment, only processors, memory, and I/O interface devices are attached to the bus, but it would alternatively be possible to attach other types of devices.

Although a specific embodiment of the invention has been disclosed along with certain alternatives, it will be recognized by those skilled in the art that additional variations in form and detail may be made within the scope of the following claims:

What is claimed is:

1. A computer system, comprising:

a plurality of devices for transmitting bus commands;

a central repeater for broadcasting commands;

a plurality of first uni-directional command bus links for transmitting commands from originating devices to said central repeater;

a uni-directional command bus broadcast portion for broadcasting commands from said central repeater to all devices attached to said command bus;

wherein each of said plurality of first uni-directional command bus links is less wide than said uni-directional command bus broadcast portion, whereby a plurality of bus cycles is required to transmit a command on one of said links, the same command being broadcast from said central repeater in a single bus cycle.

2. The computer system of claim 1, wherein said system comprises a plurality of processors, at least some of said devices for transmitting bus commands being processors.

3. The computer system of claim 1, further comprising:

a global command bus arbitrator, said arbitrator granting devices the right to transmit respective commands at respective bus cycles, wherein each command is broadcast by said central repeater a pre-defined number of cycles N after said arbitrator granted the right to transmit the command.

4. The computer system claim 1, further comprising:

a plurality of local request repeaters, each of said first unidirectional command bus links running from a device to an associated local request repeater; and

at least one second uni-directional command bus link, each of said at least one second links transmitting commands from a local request repeater to said central repeater.

5. The computer system claim 4, wherein said uni-directional command bus broadcast portion comprises:

a plurality of local broadcast repeaters;

at least one third uni-directional command bus link transmitting commands from said central repeater to said plurality of local broadcast repeaters; and

a plurality of fourth uni-directional command bus links, each of said fourth uni-directional command bus links transmitting commands from a respective local broadcast repeater to at least one device.

6. The computer system of claim 5, further comprising:

a global command bus arbitrator, said arbitrator granting devices the right to transmit respective commands at respective bus cycles;

wherein each command is transmitted by a device on a first uni-directional link a pre-defined number of cycles N after said arbitrator granted the right to transmit the command, is transmitted by a local request repeater on a second uni-directional link a pre-defined number of cycles M after said arbitrator granted the right to transmit the command, is transmitted by said central repeater on said at least one third link a pre-defined number of cycles L after said arbitrator granted the right to transmit the command, and is transmitted by each of said plurality of local broadcast repeaters on each respective fourth link a pre-defined number of cycles K after said arbitrator granted the right to transmit the command, wherein $N<M<L<K$.

7. A method for communicating between devices attached to a communication bus in a computer system, said communication bus including a command bus portion, said method comprising the steps of:

requesting control of said command bus portion by a first device attached to said communication bus during a first bus cycle;

receiving a grant of control of said command portion by said first device during a second bus cycle;

transmitting command information for a bus command on a first link of said command bus portion from said first device to a repeater unit during a plurality of consecu-

tive bus cycles beginning with a third bus cycle and ending on a fourth bus cycle, said third and fourth bus cycles occurring a respective pre-defined number of cycles after said second bus cycle; and

re-transmitting said command information on a second link of said command bus portion from a repeater unit to second device attached to said communication bus during a fifth bus cycle, said fifth bus cycle occurring a pre-defined number of cycles after said second bus cycle, wherein command information transmitted on said first link during multiple cycles is transmitted on said second link in a single cycle.

8. The method of claim 7, further comprising the steps of:

requesting control of a data portion of said communications bus by said second device during a sixth bus cycle, said sixth bus cycle occurring after said fifth bus cycle, said requesting control of said data portion step being responsive to said bus command;

receiving a grant of control of said data portion by said second device during a seventh bus cycle;

transmitting data on said data portion from said second device during an eighth bus cycle, said eighth bus cycle occurring a pre-defined number of cycles after said seventh bus cycle.

9. The method of claim 8, wherein said command information includes a tag uniquely identifying said bus command, and wherein said tag is transmitted simultaneously with said data on said data portion.

10. A processor for use in a multi-processor computer system, said computer system having a communications bus for communicating among processors and memory, said communications bus having a command bus portion for transmitting commands from an originating processor to other devices attached to the communications bus, said processor comprising:

a command bus request port for requesting the right to transmit a command on said command bus portion from an arbitrator;

a command bus grant port for receiving bus grant from said arbitrator;

a command output port having a plurality N of output pins for transmitting bus commands;

a command receive port having a plurality M of input pins for receiving bus commands, wherein M>N, the number of input pins M being sufficient to receive a bus command in one bus cycle; and

interface logic for transmitting bus commands on said command output port responsive to receiving bus grant from said arbitrator, wherein a bus command is transmitted from said plurality of output pins during a plurality of consecutive bus cycles beginning with a first cycle occurring a first pre-defined number of bus cycles after said processor receives bus grant and ending with a second bus cycle occurring a second pre-defined number of bus cycles after said processor receives bus grant.

* * * * *

11. A memory system comprising:

a memory controller having an interface that includes a plurality of memory subsystem ports;

a first memory subsystem including:

a buffer device having a first port and a second port, and

a plurality of memory devices coupled to the buffer device via the second port, wherein data is transferred between at least one memory device of the plurality of memory devices and the memory controller via the buffer device; and

a plurality of point-to-point links, each point-to-point link of the plurality of point-to-point links having a connection to a respective memory subsystem port of the plurality of memory subsystem ports, the plurality of point-to-point links including a first point-to-point link to connect the first port to a first memory subsystem port of the plurality of memory subsystem ports.

12. The memory system of claim 11, further including:

a plurality of connectors, wherein each connector of the plurality of connectors is connected to a respective point-to-point link of the plurality of point-to-point links; and

a plurality of memory subsystems, and wherein each memory subsystem of the plurality of memory subsystems includes:

a buffer device having a first port and a second port, wherein the first port is coupled to a respective connector of the plurality of connectors; and

a plurality of memory devices coupled to the buffer device via the second port.

13. The memory system of claim 12, further including a plurality of substrates wherein each memory subsystem of the plurality of memory subsystems is disposed on a respective substrate of the plurality of substrates.

14. The memory system of claim 11, wherein the first memory subsystem further includes a plurality of channels and a plurality of memory device select lines connected between the plurality of memory devices and the second port.

15. The memory system of claim 14, wherein each channel of the plurality of channels includes a plurality of high speed terminated signal lines.

16. The memory system of claim 11, further including a second memory subsystem including:

a buffer device having a first port and a second port, wherein the first port is connected to a second point-to-point link of the plurality of point-to-point links; and

a plurality of memory devices coupled to the buffer device via the second port.

17. The memory system of claim 11, further including a module substrate having a connector interface, wherein the first memory subsystem is disposed on the module substrate, and the buffer device is electrically connected to the connector interface, wherein the buffer device further comprises a memory address interface unit and transceives data, control and address signals between the plurality of memory devices and the connector interface.

18. The memory system of claim 11, further including a first termination coupled to the first point-to-point link to terminate a first end of the first point-to-point link, wherein the first point-to-point link is a high speed link.

19. The memory system of claim 18, further including a second termination coupled to the first point-to-point link to terminate a second end of the first point-to-point link.

20. The memory device of claim 11, wherein the buffer device comprises a write buffer, coupled to the first port, to hold data to be provided to at least one memory device of the plurality of memory devices.

21. The system according to claim 11, wherein the buffer device comprises a memory data interface unit for transceiving data between at least one point-to-point link and at least one memory device.

22. The system according to claim 11, wherein the memory controller comprises a data switch for transceiving data to and from the first memory subsystem.

23. The system according to claim 13, wherein each substrate comprises a card or assembly of cards.

24. The system according to claim 11, wherein the first point-to-point link comprises a bidirectional point-to-point data bus.

25. A memory system comprising:

a controller device;

a first buffer device having a first interface and a second interface;

a second buffer device having a first interface and a second interface;

a first point-to-point link having a first connection to the controller device and a second connection to the first interface of the first buffer device;

a first plurality of memory devices connected to the second interface of the first buffer device;

a second point-to-point link having a first connection to the controller device and a second connection to the first interface of the second buffer device; and

a second plurality of memory devices connected to the second interface of the second buffer device.

26. The memory system of claim 25, wherein the first buffer device and first plurality of memory devices are disposed on a first substrate, and the second buffer device and second plurality of memory devices are disposed on a second substrate.

27. The memory system of claim 26, further including:

a first plurality of high speed signal lines to connect the first plurality of memory devices to the second interface of the first buffer device;

a second plurality of high speed signal lines to connect the second plurality of memory devices to the second interface of the second buffer device;

a first plurality of termination elements connected to the first plurality of high speed signal lines; and

a second plurality of termination elements connected to the second plurality of high speed signal lines.

28. The memory system of claim 25, further including a third point-to-point link having a connection to the controller and a fourth point-to-point link having a connection to the controller.

29. The memory system of claim 25, further including:

a first channel to connect the first plurality of memory devices to the second interface of the first buffer device;

a second channel to connect the second plurality of memory devices to the second interface of the second buffer device;

a third channel connected to the second interface of the first buffer device;

a third plurality of memory devices electrically coupled to the third channel;

a fourth channel connected to the second interface of the second buffer device; and

a fourth plurality of memory devices electrically coupled to the fourth channel.

30. The memory system of claim 25, further including at least one termination element electrically connected to the first point-to-point link, wherein the first point-to-point link is a high speed link.

31. The memory system of claim 25, further including a module substrate having a connector interface, wherein the first buffer device is disposed on the module substrate, and the first buffer device is electrically connected to the connector interface, and wherein the first buffer device further comprises a memory address interface unit and transceives data, control and address information between the first plurality of memory devices and the connector interface.

32. The memory device of claim 25, wherein the first buffer device further includes a write buffer, coupled to the first interface of the first buffer device, to hold data to be provided to at least one memory device of the first plurality of memory devices.

33. A memory system comprising:

a controller device;

a first and second plurality of buffer devices, each buffer device of the first and second plurality of buffer devices having an interface connected to a respective plurality of memory devices;

a first and second repeater device;

a first point-to-point link having a first connection to the controller device and a second connection to the first repeater device;

a second point-to-point link having a first connection to the controller device and a second connection to the second repeater device;

a first plurality of repeater links, each repeater link having first connection to a respective buffer device of the first plurality of buffer devices, and a second connection to the first repeater device; and

a second plurality of repeater links, each repeater link having first connection to a respective buffer device of the second plurality of buffer devices and a second connection to the second repeater device.

34. The memory system of claim 33, wherein a first buffer device of the first plurality of buffer devices and a first buffer device of the second plurality of buffer devices, and corresponding pluralities of memory devices, are disposed on one of a plurality of respective module substrates.

35. The memory system of claim 33, further including a third point-to-point link having an end connected to the controller device and a fourth point-to-point link having an end connected to the controller device.

36. The memory device of claim 33, wherein each buffer device of the first plurality of buffer devices comprises a write buffer to hold data to be provided to at least one memory device of the respective plurality of memory devices.

37. The system according to claim 33, wherein each buffer device of the second plurality of buffer devices comprises a memory address interface unit for transceiving address information to and from at least one memory device.

38. The system according to claim 33, wherein the second repeater device comprises an address repeater unit for transceiving address information to and from the controller.

39. The system according to claim 38, wherein the second repeater device further comprises an address switch for transceiving the address information to and from the address repeater unit.

40. The system according to claim 33, wherein the second plurality of repeater links comprise an address bus for transmitting address information from the second repeater unit to the second plurality of buffer devices, the address information being associated with data transceived between the first repeater device and the first plurality of buffer devices.

41. A memory system comprising:

a controller device having an interface;

a first connector, second connector, and third connector;

a first point-to point link having a first connection to the interface and a second connection to the first connector;

a second point-to-point link having a first connection to the interface and a second connection to the second connector;

a third point-to-point link having a first connection to the interface and a second connection to the third connector; and

a first memory subsystem including:

a buffer device connected to the first connector; and

a plurality of memory devices connected to the buffer device, wherein at least one memory device of the plurality of memory devices transfers data to the controller device via the buffer device.

42. The memory system of claim 41, wherein the second and third connectors support coupling to a plurality of memory subsystems.

43. The memory system of claim 41, further including a module substrate having a connector interface, wherein the first memory subsystem is disposed on the module substrate, and the buffer device is electrically connected to the connector interface, and wherein the buffer device further comprises a memory address interface unit and transceives data, control and address signals between the plurality of memory devices and the connector interface.

44. The memory system of claim 41, further including a first termination to terminate the second connection of the first point-to-point link, wherein the first point-to-point link is a high speed link.

45. The memory system of claim 44, further including a second termination to terminate the first connection of the first point-to-point link.

46. The memory device of claim 41, wherein the buffer device comprises a write buffer to hold data to be provided to at least one memory device of the plurality of memory devices.

47. A memory system comprising:

a controller device;

a plurality of buffer devices, each buffer device of the plurality of buffer devices having an interface connected to a respective plurality of memory devices;

a repeater device;

a point-to-point link having a first connection to the controller device and a second connection to the repeater device;

a plurality of repeater links, each repeater link having first connection to a respective buffer device of the plurality of buffer devices, and a second connection to the repeater device.

48. A memory system comprising:

a memory controller having an interface that includes a plurality of memory subsystem ports;

a first memory subsystem including:

a memory data interface unit having a first port and a second port, and

a plurality of memory devices coupled to the memory data interface unit via the second port, wherein data is transferred between at least one memory device of the plurality of memory devices and the memory controller via the memory data interface unit; and

a plurality of point-to-point links, each point-to-point link of the plurality of point-to-point links having a connection to a respective memory subsystem port of the plurality of memory subsystem ports, the plurality of point-to-point links including a first point-to-point link to connect the first port to a first memory subsystem port of the plurality of memory subsystem ports.

49. A memory system comprising:

a controller device;

a first memory data interface unit having a first interface and a second interface;

a second memory data interface unit having a first interface and a second interface;

a first point-to-point link having a first connection to the controller device and a second connection to the first interface of the first memory data interface unit;

a first plurality of memory devices connected to the second interface of the first memory data interface unit;

a second point-to-point link having a first connection to the controller device and a second connection to the first interface of the second memory data interface unit; and

a second plurality of memory devices connected to the second interface of the second memory data interface unit.

50. A memory system comprising:

a controller device;

a plurality of memory data interface units, each memory data interface unit of the plurality of memory data interface units having an interface connected to a respective plurality of memory devices;

a repeater device;

a point-to-point link having a first connection to the controller device and a second connection to the repeater device;

a plurality of repeater links, each repeater link having first connection to a respective memory data interface unit of the plurality of memory data interface units, and a second connection to the repeater device.

51. A memory system comprising:

a controller device having an interface;

a first connector, second connector, and third connector;

a first point-to point link having a first connection to the interface and a second connection to the first connector;

a second point-to-point link having a first connection to the interface and a second connection to the second connector;

a third point-to-point link having a first connection to the interface and a second connection to the third connector; and

a first memory subsystem including:

a memory data interface unit connected to the first connector; and

a plurality of memory devices connected to the memory data interface unit, wherein at least one memory device of the plurality of memory devices transfers data to the controller device via the memory data interface unit.

52. A memory system comprising:

a data switch having an interface that includes a plurality of memory subsystem ports;

a first memory subsystem including:

a memory data interface unit having a first port and a second port, and

a plurality of memory devices coupled to the memory data interface unit via the second port, wherein data is transferred between at least one memory device of the plurality of memory devices and the data switch via the memory data interface unit; and

a plurality of point-to-point links, each point-to-point link of the plurality of point-to-point links having a connection to a respective memory subsystem port of the plurality of memory subsystem ports, the plurality of point-to-point links including a first point-to-point link to connect the first port to a first memory subsystem port of the plurality of memory subsystem ports.

53. A memory system comprising:

a data switch;

a first memory data interface unit having a first interface and a second interface;

a second memory data interface unit having a first interface and a second interface;

a first point-to-point link having a first connection to the data switch and a second connection to the first interface of the first memory data interface unit;

a first plurality of memory devices connected to the second interface of the first memory data interface unit;

a second point-to-point link having a first connection to the data switch and a second connection to the first interface of the second memory data interface unit; and

a second plurality of memory devices connected to the second interface of the second memory data interface unit.

54. A memory system comprising:

a controller device;

a plurality of memory data interface units, each memory data interface unit of the plurality of memory data interface units having an interface connected to a respective plurality of memory devices;

a data switch;

a point-to-point link having a first connection to the controller device and a second connection to the data switch;

a plurality of repeater links, each repeater link having first connection to a respective memory data interface unit of the plurality of memory data interface units, and a second connection to the data switch.

55. A memory system comprising:

a data switch having an interface;

a first connector, second connector, and third connector;

a first point-to point link having a first connection to the interface and a second connection to the first connector;

a second point-to-point link having a first connection to the interface and a second connection to the second connector;

a third point-to-point link having a first connection to the interface and a second connection to the third connector; and

a first memory subsystem including:

a memory data interface unit connected to the first connector; and

a plurality of memory devices connected to the memory data interface unit, wherein at least one memory device of the plurality of memory devices transfers data to the data switch via the memory data interface unit.

56. A memory system comprising:

a data switch having an interface that includes a plurality of memory subsystem ports;

a first memory subsystem including:

a memory data interface unit having a first port and a second port, and

a plurality of memory devices coupled to the memory data interface unit via the second port, wherein data is transferred between at least one memory device of the plurality of memory devices and the data switch via the memory data interface unit; and

a plurality of bidirectional point-to-point data buses, each bidirectional point-to-point data bus of the plurality of bidirectional point-to-point data buses having a connection to a respective memory subsystem port of the plurality of memory subsystem ports, the plurality of bidirectional point-to-point data buses including a first bidirectional point-to-point data bus to connect the first port to a first memory subsystem port of the plurality of memory subsystem ports.

57. A memory system comprising:

a data switch;

a first memory data interface unit having a first interface and a second interface;

a second memory data interface unit having a first interface and a second interface;

a first bidirectional point-to-point data bus having a first connection to the data switch and a second connection to the first interface of the first memory data interface unit;

a first plurality of memory devices connected to the second interface of the first memory data interface unit;

a second bidirectional point-to-point data bus having a first connection to the data switch and a second connection to the first interface of the second memory data interface unit; and

a second plurality of memory devices connected to the second interface of the second memory data interface unit.

58. A memory system comprising:

a controller device;

a plurality of memory data interface units, each memory data interface unit of the plurality of memory data interface units having an interface connected to a respective plurality of memory devices;

a data switch;

a point-to-point link having a first connection to the controller device and a second connection to the data switch;

a plurality of bidirectional point-to-point data buses, each bidirectional point-to-point data bus having first connection to a respective memory data interface unit of the plurality of memory data interface units, and a second connection to the data switch.

59. A memory system comprising:

a data switch having an interface;

a first connector, second connector, and third connector;

a first bidirectional point-to point data bus having a first connection to the interface and a second connection to the first connector;

a second bidirectional point-to-point data bus having a first connection to the interface and a second connection to the second connector;

a unidirectional memory data bus line having a first connection to the interface and a second connection to the third connector; and

a first memory subsystem including:

a memory data interface unit connected to the first connector; and

a plurality of memory devices connected to the memory data interface unit, wherein at least one memory device of the plurality of memory devices transfers data to the data switch via the memory data interface unit.